EXAMINER'S AMENDMENT

1. An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

2. Authorization for this examiner's amendment was given in a telephone interview with Mr. Goepel (Reg. No. 50,851), the undersigned, on March 23, 2009 and March 24, 2009.

3. The application has been amended as follows:

IN THE SPECIFICATION:

Please replace the first Paragraph on Page 2 of the Specification under the heading "RELATED APPLICATIONS" as follows:

-- This application is filed on even date with additional applications U.S. Serial No. 10/728,470 and U.S. Serial No. 10/728,471, now U.S. Pat. No. 7,457,861 which shared much common disclosure herewith and have substantially identical specifications. Accordingly said applications U.S. Serial No. 10/728,470 and U.S. Serial No. 10/728,471 are incorporated hereinto by this reference in its entirety for consistency. --

IN THE CLAIMS:

The claims of the invention have been amended as follows:

1. (Currently Amended) A host computer system comprising: ~~for optimizing Network Interface Card (NIC) based data communications between an application and a destination through a NIC wherein said NIC and said application are on a host computer system and said destination is reachable through a port on said NIC, said system, comprising:~~

a first processor;

a first memory;

a data storage device, the data storage device having an operating system stored thereon, the operating system executable by the first processor, the operating system creating an environment in which an application runs, the application communicating with a destination;

a send work queue, stored in the first memory, the send work queue comprising at least one send work queue entry, each send work queue entry comprising at least one of a descriptor and a control-path command;

a receive work queue, stored in the first memory, the receive work queue comprising at least one receive work queue entry, each receive work queue entry comprising at least one of a descriptor and a control-path command;

a Network Interface Card (NIC), the NIC comprising a port through which the application communicates with the destination, the NIC further comprising a second memory and a second processor;

a token table, stored in the first memory and associated with ~~and for substantially each of~~ said NIC~~s on said host computer system~~, the token table ~~for~~ communicating descriptors and control-path commands from the send work queue and receive work queue to the NIC; [[,]]

a Notification Request Area (NRA), stored in ~~a memory~~ the second memory, ~~on said NIC for indicating when~~ the NRA allowing the NIC ~~should~~ to notify ~~host software~~ the operating system that a ~~said~~ descriptor has completed; ~~completes,~~

a Master Completion Queue (MCQ), the MCQ being stored in the first memory and being associated with at least one of the send work queue and the receive work queue, the MCQ comprising ~~in said memory on said host computer, having a list of completed items, said items on said list being used to store~~ the completion status of descriptors posted to the associated work queue; ~~queues which are associated with a completion queue.~~

a Memory Deregistration List (MDL), the MDL being stored in the second memory, the MDL containing memory handles to be deregistered;

an MDL Insert Kernel Agent counter running in the operating system, which indicates where in the MDL a next deregistered memory handle should be inserted;

a Memory Region Table (MRT), the MRT being stored in the first memory and maintained by the MDL Insert Kernel Agent, containing all registration information for all memory regions registered on the NIC; and

a shadow MRT, the shadow MRT being stored in the second memory, which is a copy of the MRT entries for said NIC.

2-4. (Canceled)

5. (Currently Amended) ~~The method of claim 3 wherein said message descriptor posting comprises~~ A method of optimizing Network Interface Card (NIC) based data communications between an application and a destination, the communications occurring through a NIC wherein the NIC and the application are on a host computer system and the destination is reachable through a port on the NIC, the method comprising:

receiving, by ~~a call from a user application a VIA~~ Virtual Interface Architecture Provider Library (VIPL) running on the host computer system, a call from a user application to a function ~~it wants~~, the call ~~while~~ supplying said VIPL with a descriptor; [[,]]

linking the descriptor onto an ~~the~~ appropriate ~~Work Queue,~~ work queue;

atomically incrementing a ~~Token Table's~~ Last_Inserted_Entry of a token table, and storing a representation of the result in a Local_Insert_Index of ~~said Token Table,~~ the token table;

using the descriptor to fill in a Post_Array field ~~fields~~ of ~~said Token Table~~, the token table;

determining if ~~this~~ the Post_Array field filled by the descriptor is a ~~is the~~ first unaccepted Post_Array entry and if so, writing a new-entry-posted notification to an Out Post FIFO on an I/O bridge linking the NIC and the host computer system; [[.]]

operating, on receipt of an interrupt to a Message Unit from the Out Post FIFO on the NIC, a primary Do loop of a Message Unit which copies the entries from the Post_array field of the token table and, for each copied entry that is new:

determining whether the new entry contains a descriptor, and if so, fetching memory registration information for the memory region containing the descriptor and storing that information in a shadow copy of a memory registration table on the NIC;

determining if the new entry contains an indirect descriptor reference and if so:

copying the descriptor from application program associated memory;

building an I/O Resource Block (IORB) that represents the copied descriptor; and,

calling a Virtual Interface-to-Fibre Channel (Vito) Protocol component and passing said IORB to said Vito Protocol component;

if the new entry does not contain an indirect descriptor reference, handling the entry as a control command, else incrementing a counter for a last accepted entry; and

continuing the primary Do loop until the last copied entry has been handled.

6. (Currently Amended) The method of claim 5, contents of the descriptor being validated ~~wherein~~ prior to linking said descriptor onto said appropriate work queue~~, validating said descriptor's contents~~.

7. (Currently Amended) The method of claim 5 further comprising ~~wherein prior to using the descriptor to fill in a Post_Array fields of said Token Table,~~ determining if the ~~Token Table~~ token table is full, and if so, causing the system to wait until there is room on the token table before processing the application request, the determining occurring prior to the use of the descriptor to fill in the Post_Array field of the token table.

8-9. (Canceled)

10. (Currently Amended) The method of claim [[9]] 5 further comprising: [[,]]

adjusting the number of token table entries copied per copy operation based upon load.

11. (Currently Amended) The method of claim [[9]] 5 further comprising:

adjusting frequency of polling of the ~~Token Table~~ token table by the Message Unit based on measured load as determined by number of valid Post Array entries.

12-16. (Canceled)

17. (Currently Amended) ~~A system as set forth in~~ The host computer system of claim 1 ~~wherein there exists~~ further comprising a master completion queue for each NIC, and each ~~said~~ such master completion queue ~~can reference~~ able to reference any completion queue associated with the NIC.

18. (Currently Amended) ~~A system as set forth in~~ The host computer system of claim 1 ~~claim 2~~ wherein there is a send and a receive work queue and a completion queue for each application program and wherein a notification request area exists for each completion queue and for said master completion queue, said notification request area providing notice to requestors of completions.

19-21. (Canceled)

ALLOWABLE SUBJECT MATTER

4. Claims 1, 5-7, 10-11 and 17-18 (renumbered as 1-8) are allowable over the prior art of record.

5. The following is an examiner's statements of reason for allowance:

The examiner has found that the prior art of record does not appear to teach or suggest or render obvious the claimed limitations in combination with the specific added limitations as recited in independent claims and subsequent dependent claims. The

prior art of record fails to teach or suggest a method of optimizing Network Interface Card (NIC) based data communications between an application and a destination, the communications occurring through a NIC wherein the NIC and the application are on a host computer system and the destination is reachable through a port on the NIC, comprising receiving a call from a user application to a function by Virtual Interface Architecture Provider Library (VIPL) running on the host computer system wherein the call supplying said VIPL with a descriptor and determining if a Post_Array field of a token table filled by the descriptor is a first unaccepted Post_Array entry and if so, writing a new-entry-posted notification to an Out Post FIFO on an I/O bridge linking the NIC and the host computer system, upon receipt of an interrupt to a Message Unit from the Out Post FIFO on the NIC, operating a primary Do loop of a Message Unit which copies the entries from the Post_array field of the token table and, for each copied entry that is new, then determining whether the new entry contains a descriptor, and if so, fetching memory registration information for the memory region containing the descriptor and storing that information in a shadow copy of a memory registration table on the NIC, determining if the new entry contains an indirect descriptor reference and if so: copying the descriptor from application program associated memory, building an I/O Resource Block (IORB) that represents the copied descriptor, and calling a Virtual Interface-to-Fibre Channel (Vito) Protocol component and passing said IORB to said Vito Protocol component, if the new entry does not contain an indirect descriptor reference, handling the entry as a control command, else incrementing a counter for a

last accepted entry, and continuing the primary Do loop until the last copied entry has been handled.

In addition, the prior art of record fails to teach or suggest a system of optimizing Network Interface Card (NIC) based data communications between an application and a destination, the communications occurring through a NIC wherein the NIC and the application are on a host computer system and the destination is reachable through a port on the NIC, comprising a data storage device having an operating system stored thereon executable by a processor, the operating system creating an environment in which an application runs and wherein the application communicating with a destination, and a Network Interface Card (NIC) comprising a port through which the application communicates with the destination, and a token table stored in a first memory and associated with the NIC and communicating descriptors and control-path commands from a send work queue and a receive work queue to the NIC, and a Notification Request Area (NRA) stored in a second memory of the NIC and allowing the NIC to notify the operating system that a descriptor has completed, and a Master Completion Queue (MCQ) being stored in the first memory of the host and being associated with at least one of the send work queue and the receive work queue, and comprising the completion status of descriptors posted to the associated work queue, and a Memory Deregistration List (MDL) being stored in the second memory and containing memory handles to be deregistered and an MDL Insert Kernel Agent counter running in the operating system which indicates where in the MDL a next deregistered memory handle should be inserted.

ADDITIONAL REFERENCES

6. The examiner as of general interest cites the following references:

A) Dwork et al, U.S. Pat. No. 6,963,946.

B) Andjelic, U.S. Pat. No. 7,451,456.

C) Futral et al, U.S. Pat. No. 5,991,797.

D) Hirschfeld et al, U.S. Pat. No. 6,880,002.

E) Bakshi et al, U.S. Pat. No. 6,742,051.

F) Davis, U.S. Pat. No. 7,281,030.

G) Shah et al, U.S. Pat. No. 6,460,080.

H) Rajamony et al, U.S. Pat. No. 7,089,282.

I) Burns et al, U.S. Pat. No. 6,345,301.

J) Shah et al, U.S. Pat. No. 7,039,922.

K) Terrell et al, U.S. Pat. Application Pub. No. US 2002/0124108 A1.

7. Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

8. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Philip B. Tran whose telephone number is (571) 272-3991. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

9. Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/Philip B Tran/
Primary Examiner, Art Unit 2455
Mar 25, 2009